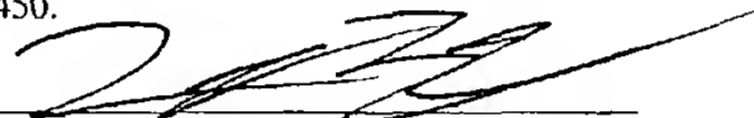


PATENT
5760-12300
VRTS 0390

"EXPRESS MAIL" MAILING LABEL NUMBER
EV 318249145 US

DATE OF DEPOSIT JULY 18, 2003

I HEREBY CERTIFY THAT THIS PAPER OR
FEE IS BEING DEPOSITED WITH THE
UNITED STATES POSTAL SERVICE
"EXPRESS MAIL POST OFFICE TO
ADDRESSEE" SERVICE UNDER 37 C.F.R. §
1.10 ON THE DATE INDICATED ABOVE AND
IS ADDRESSED TO: COMMISSIONER FOR
PATENTS, P.O. BOX 1450, ALEXANDRIA, VA
22313-1450.



Derrick Brown

Data Restore Mechanism

By:

James Philip Ohr

Robert C. Kowert
Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, Texas 78767-0398
Phone: (512) 853-8800

BACKGROUND OF THE INVENTION

Field of the Invention

5 **[0001]** This invention is related to the field of computer systems and, more particularly, to restoring data from backup storage.

Description of the Related Art

10 **[0002]** In disaster recovery backups, data is physically transferred from the primary storage media to the backup media. The backup may be to either disk or tape, though tape has traditionally dominated this market. With the continuing reduction in the cost of disk storage more sites are switching to disks as the backup media. In addition to the lower cost, disk storage tends to occupy less space and is faster than tape. While disk
15 tends to be faster than tape, it should be noted that disk backups and restores typically result in a considerable amount of application down time (typically hours).

[0003] In high-end applications, primary storage disks are typically high performance (e.g. EMC, Hitachi, or IBM arrays). Purchasing and maintaining equivalent sets of disk
20 arrays to perform mirroring can be very expensive. Therefore, many sites use inexpensive, mediocre-performance solutions for backup storage (e.g. arrays of IDE disks). Typically, users of such high-end applications do not use such backup storage as “mirrors” that can be switched to and run off backup storage due to the poor performance of the backup storage. For this and other reasons, mirroring and switching to a backup
25 image to run in a production system may not be a viable solution for many enterprises.

[0004] In addition, disaster recovery backups are typically not just copies of data like mirrors. A backup application may include backup-specific information or formatting with the backed-up data. A backup application may write to disk like it is writing to tape,
30 e.g. in TAR format. Therefore, the backed-up data in backup storage may not be in a

format that can be switched to directly to serve as the primary data in a production system.

5 [0005] In general, data moved to or from storage devices is provided using either block-level or file-level access. File level access requires some knowledge of the underlying file system and/or volume management system used to organize data on the storage devices. This type of information is typically available only at the host level, and thus I/O operations utilizing file-level access must be performed or at least managed by software executing on a host computer. Block-level access uses physical storage device addresses
10 to access data and thus need not be "assisted" by some entity having file system and/or volume knowledge.

[0006] A data restore application may restore data from backup storage to primary storage using the addresses of the source and destination devices and blocks. Such
15 address information is typically in the form of an extent list having one or more extents. An extent is typically a contiguous set of storage blocks allocated for a file portion, a file, or multiple files. Extents are typically represented by a device address indication, a starting block address on that device, and a length (number of contiguous blocks). However, extents can be defined in a variety of different ways, e.g., a starting address and
20 an ending address, no device information explicitly included, etc. Thus, an extent is generally any information used to locate a desired portion of a storage resource.

[0007] Typically, during restores, an application will have to wait for a file to be fully restored before accessing the file. Since a restore operation may restore files in any order,
25 an application may have to wait a considerable amount of time for a particular file to be fully restored. Large databases may include hundreds of gigabytes or even terabytes of data; restores of these databases may take hours or even days before the data reaches a stable state. In many cases, applications may have to wait until all of the data is restored before they can access any of the data.

30

[0008] Therefore, it is desirable to provide a restore mechanism that has reduced impact on production applications. It is also desirable to restore data needed from disk-based disaster recovery backups in a near instantaneous manner from the production application's perspective. It is also desirable to allow application to be active and
5 accessing data being restored while the restore is in progress transparent to the applications.

SUMMARY

[0009] Embodiments of a system and method for performing restores from backups while applications are active and accessing the data being restored are described. Embodiments may provide a restore mechanism that may restore data in near real-time from a disk-based backup through a coupling of a restore application with a file system and/or volume manager. Embodiments may allow restoring data into a file system while one or more applications that may use the data being restored are active. Embodiments may allow users to get backup data from backup storage onto primary storage as rapidly as possible while the restore is taking place with limited or no impact on the application(s).

[0010] To perform restores from backups while applications are active and accessing the data being restored according to one embodiment, a map correlating destination locations on primary storage to source locations on backup storage for a set of files to be restored may be generated. A restore of the set of files from the backup storage to the primary storage may be started. During the restore, it may be determined that one or more blocks of data of a file in the set of files is needed by an application. The map may be accessed to determine if the blocks have been restored. If the blocks have not been restored, the blocks may be restored from the backup storage to the primary storage. The restored blocks of data are accessible by the application while the restore is in progress. The map may be updated to indicate blocks of data that have been restored to the primary storage.

[0011] One embodiment may generate a map that depicts the file system after the restore is complete. The map describes the blocks that will be restored to the file system and their origin location on the backup disk. By providing the map of the blocks being restored to the file system, and continuously updating the map as to which block have been restored, the file system may determine if a block has already been restored, and if not to request an immediate restore of any required blocks.

[0012] Embodiments may be implemented in Storage Area Network (SAN) environments or other types of network storage environments. Embodiments may also be implemented in non-networked storage environments, for example in a single-machine system.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] The following detailed description makes reference to the accompanying drawings, which are now briefly described.

5

[0014] Figure 1 illustrates a network environment in which the restore mechanism may be implemented according to one embodiment.

[0015] Figure 2 illustrates the restore mechanism with a map correlating source locations
10 to destination locations for data of a restore according to one embodiment.

[0016] Figure 3 illustrates a map that correlates source locations to destination locations for data of a restore according to one embodiment.

15 [0017] Figure 4 illustrates the restore mechanism in an environment with a media server according to one embodiment.

[0018] Figure 5 is a flowchart of a method for performing restores from backups while applications are active and accessing the data being restored according to one
20 embodiment.

[0019] While the invention is described herein by way of example for several embodiments and illustrative drawings, those skilled in the art will recognize that the invention is not limited to the embodiments or drawings described. It should be
25 understood, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims. The headings used herein are for organizational purposes only and are not meant to be used to limit the scope of the
30 description or the claims. As used throughout this application, the word "may" is used in

a permissive sense (i.e., meaning having the potential to), rather than the mandatory sense (i.e., meaning must). Similarly, the words “include”, “including”, and “includes” mean including, but not limited to.

DETAILED DESCRIPTION OF EMBODIMENTS

[0020] Embodiments of a system and method for performing restores from backups while applications are active and accessing the data being restored are described. Embodiments
5 may provide a restore mechanism that may restore data in near real-time from a disk-based backup through a coupling of the restore application with the file system and/or volume manager. Using embodiments, a block-level restore may be performed while the application(s) that access the data is active. In one embodiment, the file system and/or volume manager may determine if blocks of data needed by active applications have been
10 restored. In one embodiment, if a block has not yet been restored, the file system and/or volume manager generates a request to the restore application to have the block immediately restored. In another embodiment, the file system or volume manager may have direct access to the backup storage, and thus may directly access and restore the needed block(s) without going through the restore application. In this embodiment, the
15 restore application may determine or be notified that the needed blocks have been restored to avoid overwriting the blocks. Embodiments of the restore mechanism may allow blocks of data to be restored to primary storage on demand and out-of-order from backup storage.

20 [0021] Embodiments may allow restoring data into a file system while one or more applications that may use the data being restored are active. In the prior art, an application may have to wait for a file to be fully restored before accessing the file, and in many cases applications may have to wait until all of the data is restored before they can access the data. Embodiments may allow users to get backup data from backup storage
25 onto primary storage as rapidly as possible while the restore is taking place with limited or no impact on the application(s).

[0022] One embodiment may generate a map that depicts the file system after the restore is complete. In one embodiment, this map may be generated by the restore application. In
30 another embodiment, the map may be generated by the file system. The map describes

the blocks that will be restored to the file system and their origin location on the backup disk. By providing the map of the blocks being restored to the file system, and continuously updating the map as to which block have been restored, the file system may determine if a block has already been restored, and if not to request an immediate restore of any required blocks.

[0023] Embodiments may be implemented in Storage Area Network (SAN) environments or other types of network storage environments. Embodiments may also be implemented in non-networked storage environments, even in a single-machine system.

[0024] Figure 1 illustrates a network environment in which the restore mechanism may be implemented according to one embodiment. In one embodiment, the restore mechanism may map the files to be restored. In one embodiment, file system 110 and restore application 112 may be on different servers (e.g. in this illustration, file system 110 is on file server 102, and restore application 112 may be on another server such as a media server). In one embodiment, restore application 112 may be on file server 102 with file system 110. In one embodiment, part of restore application 112 may be on file server 102 (e.g. an engine or driver) while the rest of restore application may be on another server (e.g. a media server).

[0025] In one embodiment, the file system 110 may allocate blocks for the files on the primary (destination) storage 114; there may be no application data in the blocks when allocated. The locations on the backup storage 116 where the data to be restored is located may be determined. In one embodiment, the restore application 112 may perform the determination. A correspondence or map of where the data is coming from on the backup storage and where the data is going to on the primary storage may be generated. This map may be a bitmap, linked list, or any other suitable data structure. This map may pair the source and destination of blocks to be restored. The map may include indications of whether particular blocks or extents including blocks have been restored. This map may be located on the file server 102, a media server (not shown), on a server with the

restore application 112, or on any server (or storage system) in the network environment where it is accessible by both the file system 110 and the restore application 112.

5 [0026] In one embodiment, the map generation may be performed by coupling the restore application 112 with the file system 110. In one embodiment, the generation of the map may be performed at the file server level. In one embodiment, the file system 110 may generate the map. During the restore operation, the map may be updated to indicate which blocks or extents have been restored. In one embodiment, the restore application 112 may maintain the map. In another embodiment, the file system 110 may maintain the
10 map. In yet another embodiment, both the file system 110 and the restore application 112 may maintain the map. In other embodiments, other entities (e.g. a driver or engine on file server 102 or a volume manager) may maintain the map.

[0027] One embodiment may include a media server (not shown) that may perform block
15 detection (e.g. detecting when blocks are needed by applications running during the restore), checking of the map to determine if needed blocks have been restored, and on-demand requesting of non-restored blocks. When the file system 110 or alternatively a media server determines that an application needs a block of data in the restore, the file system 110 or media server may examine the map to determine if the block has been
20 restored. In one embodiment, during the restore, when an application needs a block that has not been restored from the backup storage 116 to the primary storage 114 by the restore application 112, the file system 110 or alternatively a media server operating in conjunction with the file system sends a request to the restore application 112 to immediately restore the block from the backup storage 116 to the primary storage 114.
25 Alternatively, the restore application 112 may provide the requested block directly to the requestor (e.g. file system 110).

[0028] In one embodiment, detection and request for the immediate restore of non-restored blocks occurs within the file server 102. In this embodiment, detection and
30 request for blocks is performed at a software level on the file server 102 and not at the

storage hardware level. In one embodiment, detection of needed blocks may be performed by the file system 110. In this embodiment, the file system 110 determines that there is a block that is needed that has not yet been restored by examining the map and makes a request to the restore application 112 to immediately restore the block. In one embodiment, a driver on the file server 102, rather than the file system 110, may perform the detection and request for non-restored blocks. Alternatively, a media server between the file server 110 and the primary and backup storage may perform non-restored block detection and requests for the immediate restore of non-restored blocks.

10 [0029] Referring to Figure 1, in some network storage environments such as SAN environments, the file server 102 may have direct access to the backup storage 116. In these environments, in one embodiment, the file server 102 may retrieve blocks from the backup storage 116 that it needs and that have not been restored as indicated by the map. The file server 102 may then update the map to indicate that the block has been retrieved and restored. Thus, in this embodiment, the file server 102 may satisfy the on-demand request on its own without sending a request to the restore application 112. In this embodiment, the file server 102 may do the work that the restore application 112 would otherwise have to do, preferably reducing the number of messages and other operations that have to be performed and thus causing less impact to the overall restore process.

20

[0030] In this embodiment, the file server 102 may coordinate on-demand restores with the restore application 112, e.g. by updating the map, to prevent the restore application 112 from overwriting blocks restored to primary storage 114 with potentially older data from the backup storage 116. In one embodiment, the file server 102 may update the map to indicate that a retrieved block has been restored, and the restore application 112 may check the map to determine if blocks it is about to restore have already been restored and not restore any blocks that have been restored directly by the file server 102 to thus avoid overwriting blocks already retrieved by the file server 102 and possibly modified by applications. In another embodiment, the file server 102 may notify the restore

application 112 when it directly retrieves blocks and the restore application 112 may update the map.

5 [0031] One embodiment may include a media server between the file server 110 and the storage. In one embodiment, the restore application 112 may run on the media server. In one embodiment, the media server may handle detection of blocks, checking of the map, and interacting with the restore application for the on-demand restore of detected blocks. In another embodiment, the file server may handle detection of blocks, checking of the map, and interacting with the restore application for the on-demand restore of detected
10 blocks. In one embodiment, a file system on the file server may perform these tasks. In another embodiment, a driver on the file server may perform these tasks. One embodiment may not have a separate media server. In one embodiment, the restore application may run in the file server. One embodiment may be implemented on one server and one storage network. In one embodiment, the restore application may have an
15 engine running on the file server that moves the data and that is coupled with the file system; the map may be maintained in the file server. The engine may move the data from backup storage to primary storage. If file system needs some blocks that have not been restored, the file system notifies the engine to get and restore the indicated blocks.

20 [0032] One embodiment may be implemented in network environments that include a volume manager. A volume manager typically sits under the file system 110 and is used to aggregate groups of storage devices together to form larger views of storage (e.g. striped or concatenated). A volume manager may provide a uniform, singular space in which the file system 110 may operate. Multiple disks can be made to appear as one
25 storage system to the file system 110. In this embodiment, the file system 110 may perform the pre-mapping of the files. The file system 110 may allocate storage where the data is to be restored, as files may typically be identified at the file system 110 level and not at the volume manager level. In one embodiment, detection of blocks, checking of the map, and making on-demand requests to the restore application for the restore of
30 needed blocks may be performed at the file system 110 level. In another embodiment,

block detection, map checking, and on-demand requests to the restore application may be performed at the volume manager level.

5 [0033] In embodiments, address spaces may be translated; for example, the file system 110 addresses at the file level, other levels may address at the block or extent level, and lower levels address at the physical level (e.g. using LUNs). Embodiments may include or alternatively access a mapping mechanism that may be used to map addresses to whatever layer is necessary to perform the mapping and/or on-demand restore operations.

10 [0034] Figure 2 illustrates the restore mechanism with the map according to one embodiment. Figure 2 illustrates means for restoring a set of files from a backup storage to a primary storage, means for determining on a file server that one or more blocks of data of a file in the set of files needed by an application have not been restored during the restore, and means for restoring the determined one or more blocks of data according to
15 one embodiment.

[0035] In one embodiment, a set of files may need to be restored. The restore application 112 may be requested to restore the set of files. The restore application 112 may communicate with the file system 110 to inform the file system 110 to pre-allocate the set
20 of files. The restore application 112 may provide the file names and size of the files, and potentially other information about the set of files. The file system 110 may pre-allocate space (blocks or extents) for the set of files on primary storage 114 and return to the restore application 112 information describing the set of destination blocks (or extents) on the primary storage 114 to where the data is to be restored. The restore application
25 112 then may pair that set of destination blocks on the primary storage 114 with the source locations of the blocks on the backup storage 116 to generate a map 120.

[0036] Figure 3 illustrates a map according to one embodiment. In Figure 3, primary storage block information 122 is correlated with backup storage block information 124
30 for the N blocks or extents to be restored. Map 120 may be a bitmap, linked list, or any

other suitable data structure. In one embodiment, for each file in the restore, there may be a map 120 generated for that file that correlates source and destination information for the file. In another embodiment, there may be one map 120 generated that correlates source and destination information for all files in the restore. Other embodiments may generate
5 separate maps 120 for each of two or more sets of files. In one embodiment, this mapping of source and destination information may be performed for all files to be restored up front, before the restore of the files actually starts. Therefore, all the blocks for all the files to be restored may be pre-mapped at the beginning of the restore process. This pre-allocation and pre-mapping process may, for example, take seconds to minutes.

10

[0037] Referring again to Figure 2, in one embodiment, the file system 110 may maintain the map 120 that it uses to determine what blocks of the set of files that are being restored are currently valid in (restored to) the primary storage 114. The map 120 may be dynamically updated as the restore is performed, in one embodiment by the file system
15 110, in another embodiment by the restore application 112, or in yet another embodiment by both. In some embodiments, other entities such as a driver or a media server may access and/or update the map 120. In one embodiment, the restore application 112 may keep the map updated to indicate which blocks have been restored. In another embodiment, the restore application 112 may send messages to the file system 110
20 indicating which blocks have been restored, and the file system 110 may update the map 120. In one embodiment, both the file system 110 and the restore application 112 may update the map 120 when necessary. In some embodiments, other entities such as a media server or a driver on the file server may maintain and/or update the map 120.

25 [0038] In one embodiment, when the file system 110 needs to access a block on the primary storage 114, it checks the map 120 to see if the block has been restored. If the block has not been restored, the file system 110 knows that it cannot access the block directly from the primary storage 114. The file system 110 then sends a request to the restore application 112 that indicates that the file system 110 needs the block
30 immediately. The restore application 112 then goes to the backup storage 116 and gets

the block. In one embodiment, the restore application 112 may restore the requested block to the primary storage 116 and notify the file system 110 that the block has been restored. In another embodiment, the restore application 112 may provide the block directly to the file system 110, which may then write the block to the primary storage 114.

5 In both embodiments, the map 120 is updated to indicate that the block has been restored. In one embodiment, the file system 110 may update the map 120. In another embodiment, the restore application 112 may update the map 120. Alternatively, the restore application 112 may restore the block to the primary storage 114, update the map 120, and the file system 110 may check the map 120 to detect if the block has been

10 restored.

[0039] Thus, in embodiments, the restore application 112 may get the block from the backup storage 116 and make it available to the file system 110 in response to the file system 110 sending a message to the restore application 112 indicating it needs the block.

15 The map 120 is updated to indicate the block has been restored. In addition, the map 120 is updated to indicate non-requested blocks restored to primary storage 114 by the restore application 112 in the course of the normal restore process. The file system 110 may thus provide blocks of files from the backup storage 116 to an application for access (read or write) while the restore is in progress. At the same time the file system 110 may be

20 making on-demand requests for needed blocks to be immediately restored, the restore may be proceeding as normal; the restore application 112 may be moving other, non-requested blocks from the backup storage 116 to the primary storage 114. The restore may be proceeding normally in the background while on-demand restores may be occurring if the file system 110 determines blocks that it needs have not yet been restored.

25 [0040] While the data is being moved, the map 120 is being updated so that the map 120 reflects what has been restored to the primary storage 116. If the file system 110 checks the map 120 and sees that a block it needs has not yet been restored, the file system 110 notifies the restore application 112 to provide the block immediately.

30

[0041] In one embodiment, if a file access by an application does not involve a file that is being restored, then the file system 110 may determine that it does not have to check the map to determine if the file's blocks have been restored. In one embodiment, the restore application 112 may inform the file system 110 when the restore has completed so that
5 the file system 110 will know it no longer needs to check the map 120 and the map 120 may be disposed of if desired.

[0042] In one embodiment, each file's metadata may include an indication to mark if the file is to be restored. When the file system 110 receives a request for a file or a portion of
10 a file, the file system 110 may check the metadata for the file to determine if the file is to be restored. If it is to be restored, then the file system may check the map 120 for that file to determine if the needed blocks have been restored. If the needed blocks have not been restored, then the file system 110 may send a request to the restore application 112 to immediately restore the needed blocks. In one embodiment, the file system 110 may
15 check the file's metadata to determine if the file has been restored; if the file has been restored, then the file system 110 can serve the request without checking the map 120; otherwise, the file system 110 checks the map 120 and, if the map 120 indicates the needed blocks have not been restored, notifies the restore application 112 to restore the needed non-restored blocks.

20

[0043] Figure 4 illustrates the restore mechanism in an environment with a media server according to one embodiment. Primary storage 206 may be, for example, a disk array that holds the data being accessed by the file server 200 and to which a restore is being performed. The backup storage 204 may hold the data that was previously stored as part
25 of a backup operation. In one embodiment, the restore application may reside primarily on the Media Server 202. In one embodiment, some components of the restore application may reside on the File Server (client) 200. In one embodiment, when a request is made to restore some files, the files may be pre-allocated and mapped by the restore application using the capabilities of the file system on file server 200. The extents
30 pre-allocated by the file server 200 may be transferred to the Media Server 202, and the

restore application may correlate the location of the data on the backup storage 204 to the extents on the Primary Storage 206. Once the correlation is completed, the media server 202 may provide the file system on file serve 200 a map of the blocks that are in the process of being restored. The file system may use this map to make a determination if the data on the primary storage 206 is current or is in the process of being restored. If a data block is required that has not yet been restored to the primary storage 206, then the file system may make a request to the restore application for an immediate restore. The restore application may preferably immediately retrieve requested data blocks and mark the blocks as having been restored. The file system may then proceed with the storage request. The restore application may be concurrently restoring data from the backup storage 204 to the primary storage 206. The map may be updated, for example, but not necessarily, at regular intervals, to indicate the blocks that have been restored.

[0044] In one embodiment, the backup storage 206 may be directly accessible to the File Server 200. In this embodiment, the file system may directly read blocks from the backup storage 204 based on the extent mappings created by the restore application.

[0045] Embodiments may include at least some integration between the file system and the restore application. In one embodiment, the file system pre-allocates and maps the restore storage on primary storage 206. In one embodiment, the file system checks the map (e.g. bitmap, linked list or other structure) to determine if the current block being accessed has been restored. If a block has not yet been restored, then the file system requests the block be immediately restored or, alternatively, accesses the backup storage 206 for the specific block. When notified of the block's availability, the file system may proceed with the I/O access (e.g. generated by an application). In one embodiment, during the restore process with on-demand restores of blocks, the file system may run in a degraded state; however, customers may prefer to run in a degraded state than to have the application down throughout the restore.

[0046] One embodiment may include a driver under the file system to monitor the requested blocks and provide the on-demand requests. This embodiment may be used, for example, in environments where the restore application cannot be (fully) integrated with the file system. In this embodiment, pre-allocation and mapping may be performed
5 by the file system.

[0047] Figure 5 is a flowchart of a method for performing restores from backups while applications are active and accessing the data being restored according to one embodiment. As indicated at 300, a map correlating destination locations on primary
10 storage to source locations on backup storage for a set of files to be restored may be generated. A restore of the set of files from the backup storage to the primary storage may be started as indicated at 302. As indicated at 304, during the restore, it may be determined that one or more blocks of data of a file in the set of files is needed by an application. As indicated at 306, the map may be accessed to determine if the blocks
15 have been restored. As indicated at 308, if the blocks have not been restored, the blocks may be restored from the backup storage to the primary storage. The restored one or more blocks of data are accessible by the application while the restore is in progress. The map may be updated to indicate blocks of data that have been restored to the primary storage.

20

[0048] In one embodiment, a restore application performs the map generation and the restore of the set of files from the backup storage to the primary storage. In one embodiment, a file system performs the determining of the one or more blocks of data needed by the application and accessing the map to determine if the blocks have been
25 restored. In one embodiment, if it is determined that the blocks have not been restored, the file system sends a message to the restore application to instruct the restore application to restore the blocks. In this embodiment, the restore application restores the blocks of data to the primary storage in response to the message.

Conclusion

[0049] Various embodiments may further include receiving, sending or storing instructions and/or data implemented in accordance with the foregoing description upon a carrier medium. Generally speaking, a carrier medium may include storage media or memory media such as magnetic or optical media, e.g., disk or CD-ROM, volatile or non-volatile media such as RAM (e.g. SDRAM, DDR SDRAM, RDRAM, SRAM, etc.), ROM, etc. As well as transmission media or signals such as electrical, electromagnetic, or digital signals, conveyed via a communication medium such as network and/or a wireless link.

[0050] The various methods as illustrated in the Figures and described herein represent exemplary embodiments of methods. The methods may be implemented in software, hardware, or a combination thereof. The order of method may be changed, and various elements may be added, reordered, combined, omitted, modified, etc.

[0051] Various modifications and changes may be made as would be obvious to a person skilled in the art having the benefit of this disclosure. It is intended that the invention embrace all such modifications and changes and, accordingly, the above description to be regarded in an illustrative rather than a restrictive sense.